

Amendments to the Claims

Kindly amend claims 12, 14, 15 & 18, and add new claims 31-38, as set forth below. All pending claims are reproduced below, with changes in the amended claims shown by underlining (for added matter) and strikethrough/double brackets (for deleted matter).

1-11. (Previously Canceled).

12. (Currently Amended) A method of processing comprising:

providing, by a dedicated collective offload engine coupled to a switch fabric in a distributed, parallel computing system, collective processing of data from at least some processing nodes of multiple processing nodes of the distributed, parallel computing system, the collective processing of data by the dedicated collective offload engine being without use of any software tree, and the dedicated collective offload engine being a hardware device coupled to the switch fabric, the hardware device being a specialized device dedicated to providing collective processing of data received from the at least some processing nodes over the switch fabric and comprising a dispatcher built from field programmable gate arrays, and a pipelined arithmetic logic unit, the dispatcher controlling collective processing of the received data by the arithmetic logic unit, the collective processing implementing a collective operation on the received data from the at least some processing nodes without use of ~~any~~ any software tree;

producing, by the dedicated collective offload engine, a result in deterministic time based on said collective processing; and

forwarding said result across the switch fabric to at least one processing node of the multiple processing nodes.

13. (Previously Presented) The method of claim 12, wherein the collective operation is a Message Passing Interface (MPI) collective operation.

14. (Currently Amended) The method of claim 12, further comprising:

receiving and storing, at a payload memory, the data from the at least some processing nodes of the multiple processing nodes, wherein said payload memory is a component of the dedicated collective offload engine; and

retrieving and performing, at ~~the~~ arithmetic logic unit (ALU), the collective processing of data stored in the payload memory, wherein said ALU is a component of the dedicated collective offload engine and is coupled to the payload memory.

15. (Currently Amended) The method of claim 14, further comprising:

controlling the collective processing of the data from the at least some processing nodes of the multiple processing nodes, wherein said controlling is performed by ~~the~~ dispatcher of the dedicated collective offload engine coupled to the ALU, and in communication with the at least some processing nodes of the multiple processing nodes via the switch fabric; and

controlling, by the dispatcher, the sharing of the result with the at least one processing node of the multiple processing nodes.

16. (Original) The method of claim 15, further comprising:

storing, in at least one task table coupled to the dispatcher, task identification information related to the at least some processing nodes of the multiple processing nodes, wherein said at least one task table is a component of the dedicated collective offload engine; and

storing, in at least one synchronization group table coupled to the dispatcher, identification information related to one or more groups of the at least some processing nodes of the multiple processing nodes, wherein said at least one synchronization group table is a component of the dedicated collective offload engine.

17. (Original) The method of claim 15, further comprising:

communicating, via an adapter, across the switch fabric using a link protocol, wherein said adapter is coupled to the switch fabric and is a component of the dedicated collective offload engine; and

facilitating, by interface logic, communication between said adapter and said payload memory and between said adapter and said dispatcher, wherein said interface logic is a component of the dedicated collective offload engine.

18. (Currently Amended) The method of claim 12, further comprising:

communicating ~~among~~ by a plurality of cascaded, dedicated collective offload engines with the at least some processing nodes via the switch fabric, wherein said communicating facilitates the collective processing of data from the at least some processing nodes of the multiple processing nodes and the producing of the result based thereon by the plurality of cascaded dedicated collective offload engines.

19. (Original) The method of claim 12, further comprising:

communicating among a plurality of dedicated collective offload engines via a channel disposed therebetween, said channel being independent of the switch fabric, wherein said communicating facilitates the collective processing of data from the at least some processing nodes of the multiple processing nodes and the producing of the result based thereon.

20-30. (Previously Canceled).

31. (New) A processing system comprising:

wherein the dedicated collective offload engine collectively processes data distributed, parallel computing system without use of any software tree, the dedicated collective offload engine being a hardware device coupled to the switch fabric, the hardware device being a specialized device dedicated to providing collective processing of data received from the at least some processing nodes across the switch fabric and comprising a dispatcher built from field programmable gate arrays, and a pipelined arithmetic logic unit, the dispatcher from at least some processing nodes of multiple processing nodes of the controlling collective processing of the received data by the arithmetic logic unit, the collective processing implementing a collective operation on the received data from the at least some processing nodes without use of any software tree;

wherein the dedicated collective offload engine produces a result in deterministic time based on the collective processing; and

wherein the dedicated collective offload engine forwards the result across the switch fabric to at least one processing node of the multiple processing nodes.

32. (New) The processing system of claim 31, wherein the collective operation is a Message Passing Interface (MPI) collective operation.

33. (New) The processing system of claim 31, wherein the dedicated collective offload engine comprises:

a payload memory configured to receive and store the data from the at least some processing nodes of the multiple processing nodes; and

wherein the arithmetic logic unit (ALU) is coupled to the payload memory and is configured to retrieve and perform the collective processing of data stored in the payload memory.

34. (New) The processing system of claim 33, wherein the dedicated collective offload engine further comprises:

at least one task table coupled to the dispatcher, wherein the at least one task table is configured to store task identification information related to the at least some processing nodes of the multiple processing nodes; and

at least one synchronization group table coupled to the dispatcher, wherein the at least one synchronization group table is configured to store identification information related to one or more groups of the at least some processing nodes of the multiple processing nodes.

35. (New) The processing system of claim 34, wherein the dedicated collective offload engine further comprises:

an adapter coupled to the switch fabric, wherein said adapter is configured to communicate with the switch fabric using a link protocol; and

interface logic coupled to the adapter, the payload memory and the dispatcher, wherein the interface logic facilitates communication between said adapter and said payload memory and between said adapter and said dispatcher.

36. (New) The processing system of claim 31, wherein the processing system further comprises a plurality of cascaded, dedicated collective offload engines in communication with the at least some processing nodes via the switch fabric, wherein said communication facilitates the collective processing of data from the at least some processing nodes of the multiple processing nodes and the producing of the result based thereon.

37. (New) The processing system of claim 31, wherein the processing system further comprises a plurality of dedicated collective offload engines in communication with one another via a channel disposed therebetween, said channel being independent of the switch fabric, and wherein said communication facilitates the collective processing of data from the at least some processing nodes of the multiple processing nodes and the producing of the result based thereon.

38. (New) The processing system of claim 31, wherein the collective processing provided by the dedicated collective offload engine includes managing at least one distributed lock associated with at least one of a distributed database and a distributed file system.